

Multi-Ridge Fitting for Ring-Diagram Helioseismology

Benjamin J. Greer *, Bradley W. Hindman, and Juri Toomre

*JILA and Department of Astrophysical and Planetary Sciences,
University of Colorado, Boulder CO*

Dated: February 24, 2014

Abstract

Inferences of sub-surface flow velocities using local domain ring-diagram helioseismology depend on measuring the frequency splittings of oscillation modes seen in acoustic power spectra. Current methods for making these measurements utilize maximum-likelihood fitting techniques to match a model of modal power to the spectra. The model typically describes a single oscillation mode, and each mode in a given power spectrum is fit independently. We present a new method that produces measurements with greater reliability and accuracy by fitting multiple modes simultaneously. We demonstrate how this method permits measurements of sub-surface flows deeper into the Sun while providing higher uniformity in data coverage and velocity response closer to the limb of the solar disk. While the previous fitting method performs better for some measurements of low-phase-speed modes, we find this new method to be particularly useful for high phase-speed modes and small spatial areas.

1 Introduction

Helioseismology determines the structure and dynamics of the solar interior through analysis of seismic waves observed at the surface. Ring-diagram helioseismology ([12]; [2]; [9]) investigates sub-surface horizontal flows by measuring the direction-dependent frequency shift of oscillation modes. The Doppler shift of the frequency of an oscillation mode due to a sub-surface flow is expressed as

$$\delta\omega_n(\mathbf{k}) = \mathbf{k} \cdot \mathbf{u}_n(k), \quad (1)$$

*Electronic address: benjamin.greer@colorado.edu

where \mathbf{k} and n are the horizontal wavenumber and radial order of the oscillation mode, and $\mathbf{u}_n(k)$ is a spatial average of the horizontal velocity within the Sun,

$$\mathbf{u}_n(k) = \langle \mathbf{v}(\mathbf{r}) \rangle = \int K_n(\mathbf{r}; k) \mathbf{v}(\mathbf{r}) \, d^3r. \quad (2)$$

Here, $\mathbf{v}(\mathbf{r})$ is the true horizontal sub-surface flow velocity at any point \mathbf{r} in the Sun, and $K_n(\mathbf{r}; k)$ is the weighting function — or sensitivity kernel — associated with each mode (k, n) , which describes the spatial extent over which the true velocity is averaged to create a single frequency shift [3]. The variation of the frequency shift as a function of direction (the frequency splitting) is measured to provide an estimate of $\mathbf{u}_n(k)$.

While the interpretation of the frequency splitting is straightforward, the method of extracting it from the data is not. Traditionally, a model of the spectral power is fit to oscillation modes visible in power spectra of line-of-sight velocity observed in the photosphere. The model accounts for a frequency shift of the modal power as a function of horizontal direction, and this shift is directly translated into a velocity as shown in Equation 1. The specifics of the fitting procedure determine how well the frequency splittings are measured, as well as what other qualities of the power spectra are taken into account. There are currently two commonly used fitting procedures in the *Helioseismic and Magnetic Imager* (HMI) Ring-Diagram Pipeline [4]. The first method considered in this article is one introduced by [9], which fits a frequency-shifted Lorentzian model to individual modes. Since this method analyzes single ridges of modal power sequentially, we refer to it as the Single-Ridge Fitting method (SRF). The second method used in the HMI Pipeline, which will not be considered here, also fits modes independently [2], but uses a significantly different power model that includes asymmetries in modal power, two background terms, and a host of other wavenumber- and direction-dependent modifications.

In this article, we present a new fitting method that utilizes a model similar to that used in the SRF method, but modified to permit multiple radial orders to be fit simultaneously. The development of this Multi-Ridge Fitting (MRF) method is an attempt to improve upon the performance of the SRF method in terms of reliability and accuracy of measured frequency splittings. We present a comparison of the frequency splittings (interpreted as an average velocity) from a common data set processed with both fitting methods. We focus on the performance for measurements of modes that reach deepest into the Sun and for measurements made near the solar limb. As metrics of reliability and accuracy we consider the fit success rate, the typical random errors in the measurements, and how well each method recovers a known velocity.

In Section 2 we outline the data-processing steps taken to prepare solar observations for helioseismic mode fitting. In Section 3 we describe the SRF and MRF fitting methods and how they differ in procedure. In Section 4 we compare the performance of the two fitting methods using a common data set. In Section 5 we discuss the implications of the results in the context of improving accuracy and data coverage in ring-diagram helioseismology.

2 Data Acquisition

The observable we use to determine helioseismic frequencies is the line-of-sight Doppler velocity measured in the photosphere. The Doppler-velocity measurements are produced by HMI ([17]) aboard the Solar Dynamics Observatory (SDO). HMI captures full-disk Dopplergram images with a cadence of 45 seconds and an array size of 4096×4096 pixels, resulting in a spatial resolution of roughly 350 km per pixel at disk center. Tiles of various sizes (*e.g.*, 16° , 4° , 2° in heliographic angle) are extracted from the full Dopplergram images with tile centers spaced evenly on a latitude – longitude grid so that adjacent tiles overlap by half their diameter. The tile centers are shifted in longitude as time progresses at a rate corresponding to the measured surface differential rotation rate from [19]. The purpose of this tracking is to minimize the signal of differential rotation and produce a more isotropic flow field. Each tile in the mosaic of tiles for a given tile size is tracked through 25.6 hours (2048 time steps), corresponding to about 15° of solar rotation. In this article we refer to this tracking duration as one day or one realization.

The tiles are apodized in space and time and Fourier transformed to create a 3-D power spectrum for each tile. The spectra are transformed from Cartesian coordinates (k_x, k_y, ν) to polar coordinates (k, θ, ν) so that oscillation modes at a constant wavenumber k can be easily extracted. The positive x -direction ($\theta = 0$) is taken to be in the prograde zonal direction. The MRF code uses these spectra as input, while the SRF code requires an additional processing step. For the SRF method, the power is normalized by the average power computed at each wavenumber k and azimuth angle θ [10]. The purpose of this filtering is to eliminate large power variations as a function of azimuth which arise from a variety of sources including camera astigmatism and power foreshortening. The MRF method does not use filtered power spectra, as it allows for such variations of power in the fitting procedure itself.

3 Fitting Methods

3.1 Single-Ridge Fitting (SRF)

In a cylindrical section at constant wavenumber k , distinct modes can clearly be seen (Figure 1). Each radial order intersects the surface of constant wavenumber to create a band of power with a central frequency. Sub-surface flows cause each band to trace out a slight sinusoidal undulation in frequency (ν) as a function of azimuth (θ). In order to separate each radial order seen in the data, the SRF method extracts a small range of frequencies at a constant wavenumber centered around a specific mode. The bounds in frequency are based on a guess table that provides the initial values of the parameters that will be fit to the data. When a mode is framed both above and below in frequency by other modes, the bounds in frequency are taken as part-way between the central frequencies of the adjacent modes. For the highest and lowest modes in

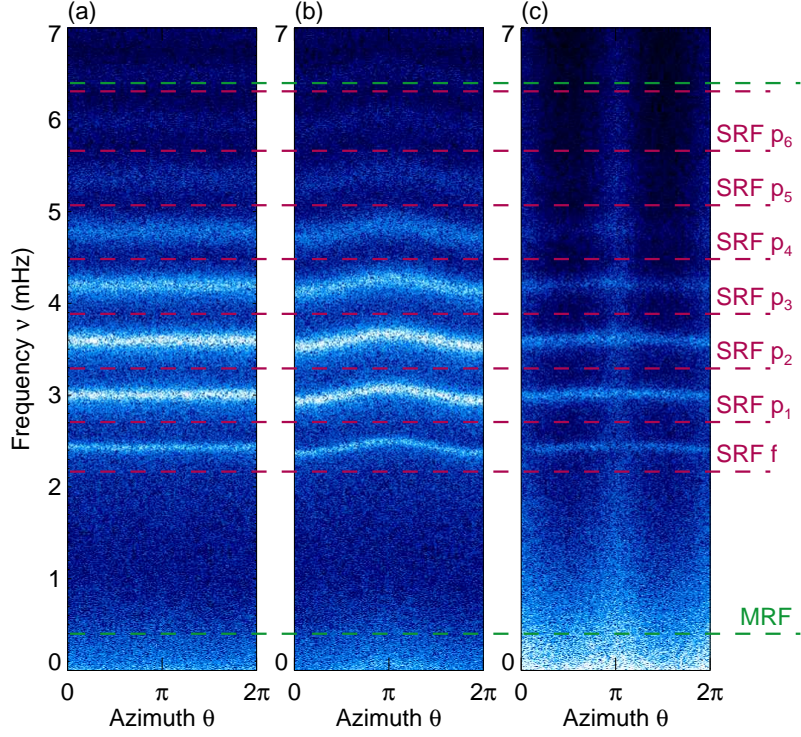


Figure 1: Sections at $k = 0.8 \text{ Mm}^{-1}$ of the power spectra for three different 16° tiles: (a) tile extracted from disk center tracked at the Snodgrass rate, (b) tile from disk center tracked at 500 m s^{-1} relative to the Snodgrass rate, and (c) tile from the central meridian and 75°N latitude tracked at 130 m s^{-1} . Red dashed lines indicate approximate bounds for the frequency windows used for independent mode fitting in the SRF method. Green dashed lines indicate the single frequency window used in the MRF method. Velocity-induced frequency shifts are evident in (b) and (c), and significant power foreshortening is visible in (c).

frequency that do not have adjacent modes on both sides, the outer boundaries are taken as a few linewidths away. SRF sequentially extracts frequency windows around each desired radial order and fits a six-parameter model (Equation 3) to each window using a maximum-likelihood technique ([1]; [9]). The SRF method represents the power in each mode as a symmetric Lorentzian with an angle-dependent frequency:

$$P(\nu, \theta) = B + \frac{A\Gamma/2}{[\nu - \nu_0 + (2\pi)^{-1}\mathbf{k} \cdot \mathbf{u}]^2 + (\Gamma/2)^2}, \quad (3)$$

$$\mathbf{u} = u_x \cos\theta \hat{\mathbf{x}} + u_y \sin\theta \hat{\mathbf{y}}. \quad (4)$$

Here $P(\nu, \theta)$ is the power as a function of frequency and azimuth; A is the amplitude; Γ is the line-width; ν_0 is the central frequency; k is the wavenumber; u_x and u_y are the Doppler velocity components; and B is a constant background. There are six parameters in this model ($A, \Gamma, \nu_0, u_x, u_y, B$) that need to be fit to each mode independently at every wavenumber. By marching through each discrete wavenumber for a given tile size and attempting to fit each mode listed in the guess table, the SRF method produces the optimal values of these six parameters for every mode. A measurement of the random error of each velocity component (u_x and u_y) is determined by the curvature of the maximum-likelihood function evaluated at the point of optimization [1]. It is this that we identify as the error or uncertainty in each measurement that we will be making.

There are a number of ways in which a fit can fail or be deemed invalid. The numerical optimization procedure can fail to converge on a solution, causing no valid data to be produced for a single mode. If there is a successful fit, the data can still be rejected if the parameters are outside of predetermined bounds. These bounds are chosen to ensure that the parameters obtained through fitting are physically relevant.

At moderate wavenumbers ($k \geq 0.5 \text{ Mm}^{-1}$) in larger tile sizes, individual modes are sufficiently separated in frequency to justify this approach of independent fitting (Figure 2). However, at lower wavenumbers, modes of neighboring radial order blend together to form a mound of spectral power without prominent individual peaks. Smaller tile sizes also have more blending of modes at all wavenumbers due to the smaller spatial apodization. In these situations, severe mode blending brings into question the idea of fitting each mode independently with a single Lorentzian model with a constant background term. The frequency window used around each mode also raises issues when a sub-surface flow introduces a large frequency splitting. At $k = 0.5 \text{ Mm}^{-1}$, a 300 m s^{-1} flow is sufficient to shift significant mode power past the frequency boundaries chosen to split up the data for fitting. Tiles produced in the standard HMI Ring-Diagram Pipeline are tracked at the Carrington rate [4], which at extreme latitudes differs from the local surface rotation rate by nearly 250 m s^{-1} . Not only does mode power shift outside the fitting window for a given mode, but power from neighboring modes that is not accounted for in the model enters the window and potentially throws off the fitting [16].

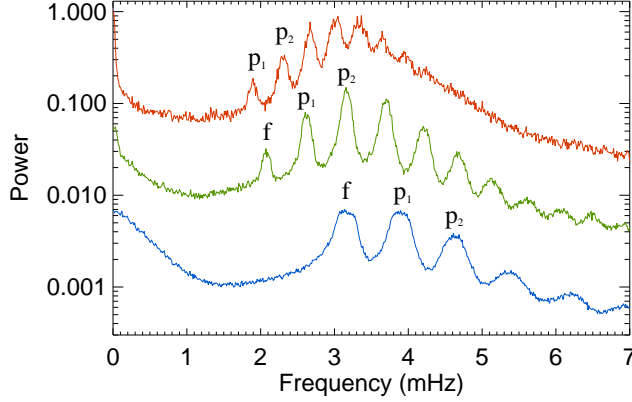


Figure 2: Azimuthal averages of a spectrum generated from a single 16° tile extracted at disk center. Each curve corresponds to a different horizontal wavenumber: $k = 0.23 \text{ Mm}^{-1}$ (red, top), $k = 0.58 \text{ Mm}^{-1}$ (green, middle), and $k = 1.42 \text{ Mm}^{-1}$ (blue, bottom). Curves are offset in power for clarity. Modes at lower wavenumber and higher frequency are spaced closer together, causing significant overlap. Ridges at high wavenumber look flattened due to azimuthal averaging over a flow-induced frequency splitting.

3.2 Multi-Ridge Fitting (MRF)

To mitigate the issues that come with independent fitting of modes, we have developed the Multi-Ridge Fitting (MRF) code. This new method fits multiple radial orders simultaneously at a given wavenumber over a wider window in frequency. The function used to describe the collection of modes is a sum of symmetric Lorentzians, each similar to those used in the SRF method:

$$P(\nu, \theta) = B(\nu, \theta) + \sum_{n=0}^{N-1} \frac{A_n(\Gamma_n/2)}{[\nu - \nu_n + (2\pi)^{-1} \mathbf{k} \cdot \mathbf{u}_n]^2 + (\Gamma_n/2)^2} F(\theta; f_n, \theta_n), \quad (5)$$

$$F(\theta; f_n, \theta_n) = 1 + f_n \cos[2(\theta - \theta_n)], \quad (6)$$

where A_n , Γ_n , ν_n , $u_{x,n}$, and $u_{y,n}$ retain the same purpose as in the SRF method, and the sum is taken over the number of modes specified in the guess table at a given wavenumber. The new factor $F(\theta; f_n, \theta_n)$ in the numerator of each Lorentzian accounts for amplitude variation as a function of θ using the parameters f_n and θ_n . While the constant background term used in Equation 3 is roughly valid when considering the narrow range of frequencies used in SRF, the background term is modified here to represent background power over a wider range by using a modified Harvey law [11]:

$$B(\nu, \theta) = \frac{B_0}{1 + (\nu/\nu_{\text{bg}})^b} F(\theta; f_{\text{bg}}, \theta_{\text{bg}}). \quad (7)$$

Here, B_0 is the amplitude; ν_{bg} is a roll-off frequency; and b is the power law index. Again, there is a power anisotropy term ($f_{\text{bg}}, \theta_{\text{bg}}$) in the amplitude. The low-frequency end of the fitting window is taken as 0.3 mHz at all wavenumbers to allow a large fraction of the background power to be utilized in constraining the model. The top of the frequency window is taken as two linewidths above the highest central frequency listed in the guess table at a given wavenumber. As a pre-conditioning step before performing the optimization of every parameter in Equation 5, the MRF method performs a simple three-parameter fit (A, Γ, ν_0) to each mode listed in the guess table. This step is similar to the fit performed in the SRF method, but does not take any variation along azimuth into account. This step provides an improved initial condition for the full optimization for each mode. Improved initial parameters for the background term are also obtained prior to the final optimization by fitting the parameters (B_0, ν_{bg}, b) to the spectrum between 0.3 mHz and 1.5 mHz. Random-error estimates are once again measured from the curvature of the maximum-likelihood function.

It is important to note the significant increase in the number of parameters from Equation 3 to Equation 5. Typical nonlinear optimization algorithms scale in time as m^2 , where m is the number of parameters being fit simultaneously. To fit all of the ridges at a single wavenumber, the SRF method performs N sequential optimizations with $m = 6$. The MRF method fits the same set of ridges simultaneously with $m = 7N + 5$, leading to an expected optimization time $\approx 1.4N$ times slower for large N . Not only is this new procedure more computationally demanding, it is much more susceptible to the numerical issue of wandering through parameter space. To ensure a timely convergence, constraints are placed on some of the model parameters:

- Amplitudes and widths must be positive.
- Central frequencies must not cross those of adjacent modes.
- Widths must be within an order of magnitude of lookup table values.
- Both zonal and meridional velocities must be smaller than $\pm 1\text{km s}^{-1}$.
- Fractional anisotropy must be between 0 and 1.

These constraints are enforced during the optimization process to limit the available parameter space.

4 Comparison

Each fitting method requires a set of initial values for each parameter, often referred to as a guess table. The guess table contains typical values of the central frequencies, widths, and amplitudes of each mode that are fit. To have a fair comparison between the two methods, they have been supplied the same guess table, and thus the same set of modes to attempt for each tile size (Figure 3).

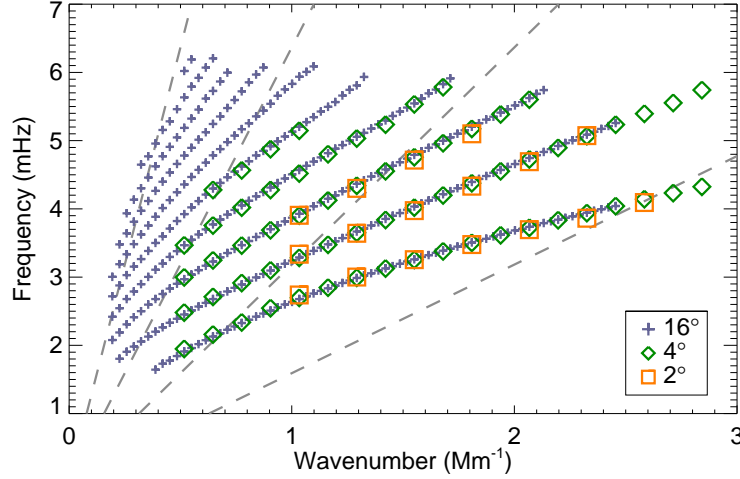


Figure 3: Guess-table frequencies for each tile size demonstrating which modes are attempted in both fitting methods. The gray dashed lines are lines of constant phase-speed (ω/k). From left to right, the values of phase-speed are 10, 20, 40, and 80 km s^{-1} . The discrete ridges of modes indicate the individual radial orders, with the f -mode at low frequency, followed by p_1 through p_9 .

While the MRF method accounts for power anisotropy in the model (Equation 5), the SRF method depends on additional processing of the input power spectra to flatten power anisotropy. The unfiltered power spectra are fit using both methods, while the filtered ones are only fit by the SRF method. We confirm that the SRF method performs slightly better on the filtered spectra (Haber, personal communication, 2013), so the analysis in this article compares unfiltered spectra fit with the MRF method to filtered spectra fit with the SRF method.

4.1 High Phase-Speed Modes

Each distinct mode that is fit has an associated kernel that specifies the spatial sensitivity to sub-surface flows (see Equation 2). While the horizontal profile of these kernels is largely determined by the tile apodization ([13]; [3]), the vertical profile changes significantly across the mode set as it depends on the mode's radial eigenfunction. The sensitivity for f -mode kernels has a single peak in depth, and each successively higher radial order adds an additional peak. The final peak of sensitivity for any mode occurs approximately where the horizontal phase-speed ω/k equals the local sound speed, causing each mode within a single radial order but different horizontal wavenumber to have a slightly different kernel. By matching phase-speeds to the sound speeds tabulated in the tabulated solar Model S [5], we can use phase-speed as a proxy for the depth to which a mode reaches into the Sun.

Current results that utilize the SRF method can reach down to 20 Mm ($0.97 R_\odot$),

partway into the near-surface shear layer [10]. In order to extend modern ring-diagram analysis deeper into the Sun, we must obtain frequency-splitting measurements for higher phase-speed modes. An easy way to do this is to increase the tile size or the tracking duration, with the result of increasing the signal-to-noise across the entire power spectrum of each tile. Increased tile sizes also suffer less blending and can sample lower wavenumbers. However, if we wish to preserve horizontal and temporal resolution, we must consider how improvements to the fitting method can produce reliable measurements for higher phase-speed modes of a constant tile size.

The first aspect to consider is how reliable each method is for obtaining measurements of high phase-speed modes. By dividing the number of successful measurements by the number attempted, we get a measure for each mode of the success rate on a per-mode basis. There are two primary reasons that either fitting method will fail to produce a successful measurement: poor data quality and poor choice of model. In the first case, as the signal-to-noise ratio of a particular mode decreases, the measured error on each model parameter increases. At some point, the fitting methods will fail to locate the mode amongst the noise and the fit will be deemed a failure. In the second case, if there is no way for the fitting method to adjust the provided model in a way that appropriately matches the data, it will struggle to converge on an optimal and unique solution. Both fitting methods will judge the attempt a failure if it does not converge within a specified number of optimization iterations or converges to non-physical model parameters. Measurements of the error on the model parameters for this case are not necessarily an accurate estimation due to the irreparable disparity between model and data.

Plotted as a function of lower turning point depth, Figure 4 demonstrates how deep into the Sun each fitting method is capable of measuring. To isolate the effects of phase-speed, the success rate illustrated here was compiled using only tiles bounded within 30° of disk center. Since the spacing between adjacent tiles in our mosaic depends on the tile size considered, there is a different number of tiles within this boundary for each tile size (49 for 16° tiles, 805 for 4° , and 3249 for 2°). We average the success rate over ten independent realizations of all tiles found within this boundary, resulting in 490 attempted velocity measurements of each mode for 16° tiles, 8050 for 4° tiles, and 32 490 for 2° tiles. While the guess table has been tuned to provide optimal results for the SRF method, the scatter seen in the SRF success rate is primarily caused by imperfect guess parameters. The MRF method is in general less susceptible to this effect. A striking difference between the two methods is how the success rate falls off with increasing depth. Both methods are most successful for the shallowest depths sampled with a given tile size. The SRF method tends to have a gradual decrease in the success rate as the depth increases, while the MRF method maintains a nearly constant success rate through the entire mode set. As the tile size decreases, the MRF success rate again stays largely constant while the SRF success rate decreases for all depths. The success rate for the SRF method also has a slight dependence on wavenumber, causing the variation of success rate between each radial order at every depth.

The measured random errors as determined by each fitting method also show a

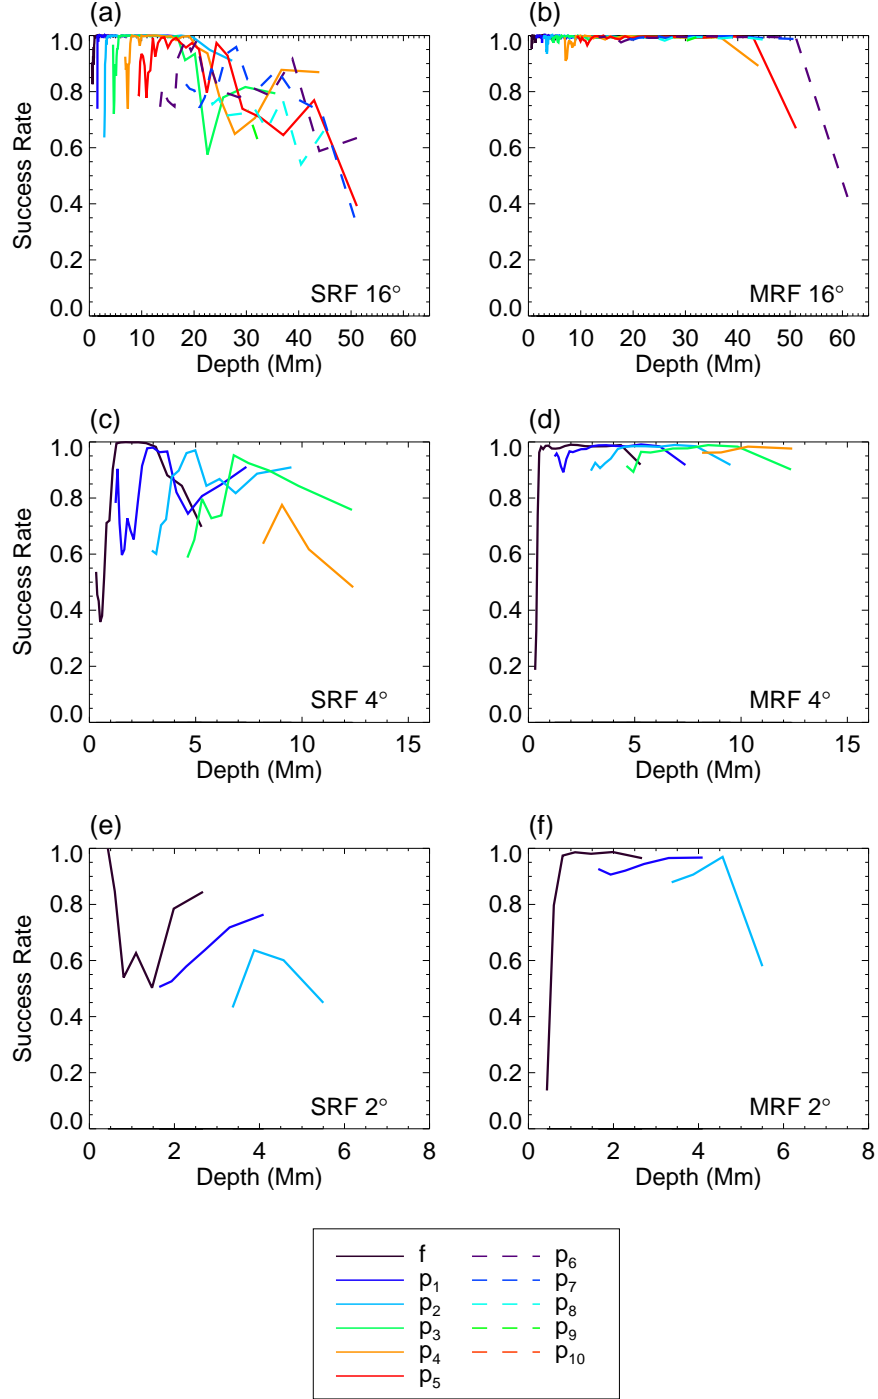


Figure 4: Fit success rate as a function of lower turning point of each mode found in the guess table. Left column is SRF results. Right column is MRF results. Rows correspond to the three tile sizes. To minimize the effects of disk position, all tiles of a given size falling within 30° of disk center have been considered here. The success rate is averaged over ten independent realizations of these tiles, covering ten days of tracking. The MRF success rate is higher at nearly every depth for every tile size.

strong dependence on phase-speed (Figure 5). As a function of lower turning-point depth for each mode, the average random error is bounded from below by an envelope that exponentially increases with depth. Near-surface flows in the upper few megameters of the Sun (typical flow speed $\approx 150 \text{ m s}^{-1}$) can be determined with sufficient precision using a single mosaic of tiles tracked through one day, while flows at the bottom of the near surface shear layer ($\approx 35 \text{ Mm}$, typical flow speed $\approx 50 \text{ m/s}$: [18]) require a significant amount of averaging in time and space to achieve a reasonable signal-to-noise ratio. Despite the considerable differences in success rate at extreme depths, there are only slight differences between the typical measured errors of each method. The average error obtained with the MRF method is slightly larger for shallow modes than that obtained with the SRF method. For deeper penetrating modes and for smaller tile sizes, the MRF errors become smaller than the SRF errors.

Through these two metrics (success rate, measured errors), we see that the MRF method performs more consistently with depth while maintaining similar estimates of the average error.

4.2 Approaching the Solar Limb

The previous results regarding how each fitting method performs at high phase-speeds were compiled using tiles near the center of the solar disk ($r \leq 30^\circ$). Both methods perform optimally near disk center and produce lower-quality results closer to the limb. The primary cause of this is foreshortening, which causes the effective resolution in one direction to be considerably less than that in the perpendicular direction. In the power spectra, this appears as a reduction of power along the direction of foreshortening. Since the MRF method has additional terms to account for this type of power variation in both the mode power and background power, it is expected that it will outperform the SRF method near the limb.

The ability to push closer to the limb has many benefits. Regions of strong magnetic activity are continually evolving as they pass across the solar disk from the east limb to the west limb. To maximize the amount of time that these regions are accessible with ring-diagram analysis, we must be able to measure flow velocities close to the solar limb. The consistency of velocity measurements as a function of disk position is important for studying the evolution of flows across the disk. The possibility of high-latitude meridional counter-cells has been of great interest ([15]; [14]; [7]), increasing the need for analysis methods that perform consistently between disk center and high latitudes. Careful removal of large-scale systematics has a noticeable effect on the determination of sub-surface flows ([21], [8], [20]). The accurate measurement of these systematics is also dependent on the spatial uniformity of the methods that obtain velocity measurements.

Once again, it is useful to first consider the success rate of each method as a function of disk position (Figure 6). A single mosaic of tiles tracked through one day covering the entire solar disk was analyzed to produce 223 750 attempted measurements for 16° tiles, 611 390 for 4° tiles, and 632 740 for 2° tiles. For the SRF method, not only

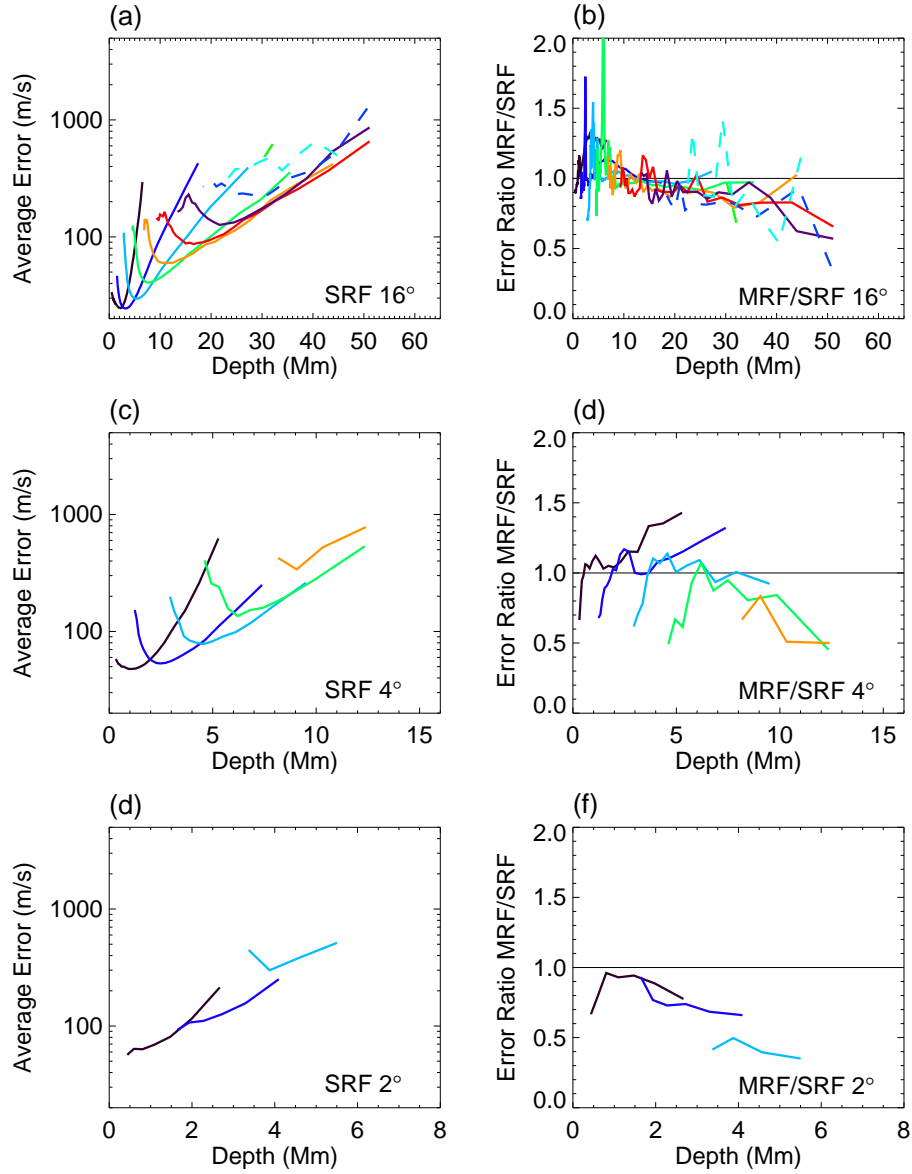


Figure 5: Left: Average random error of the SRF method as a function of lower turning point for each mode. Right: Ratio of MRF errors to SRF errors. Rows correspond to the three tile sizes. To minimize the effects of disk position, all tiles of a given size falling within 30° of disk center have been considered here. The success rate is averaged over ten independent realizations of these tiles, covering ten days of tracking. The MRF method produces errors that are largely similar to those from the SRF method, with some differences for modes that reach deepest.

does the success rate at every disk position depend on the radial order of a mode, but the success rate for each radial order also falls off at different distances from disk center. Each radial order has a distance from disk center where the success rate transitions from a nearly constant value near disk center (0°) to a steep drop towards the limb. For 16° -tile f -mode fits made with the SRF method, this distance is around 75° in heliographic angle, but for each higher order the distance decreases. By p_9 , this distance has dropped to only 45° . With fewer high phase-speed measurements being made away from disk center, the depth to which sub-surface flows can be determined becomes position dependent. The smaller tile sizes display a similar trend, only with a lower success rate for each order at all disk positions.

In contrast, the MRF method has a higher success rate in nearly all cases. For 16° tiles, the distances at which each radial order begins to drop off are now nearly equal at around 75° . This results in the MRF method producing much more uniform results across the solar disk, both in spatial coverage and in depth. Smaller tile sizes show a decrease in the success rate closer to disk center, but it is still roughly constant for all modes.

The average error as a function of distance from disk center (Figure 7) shows similar trends for both fitting methods. Both have nearly constant errors from disk center out to around 60° . Beyond this, the magnitude of the errors for all radial orders rise steeply in both fitting methods. The primary difference between the two methods is how severe this rise in error is near the limb. For 16° tiles fit with the SRF method, the average error for high-radial-order modes quickly passes 1 km s^{-1} outside of 60° from disk center. The MRF method shows a slower rise in error for these modes, reaching 1 km s^{-1} only for the worst cases.

4.3 Velocity Accuracy

As demonstrated in the previous sections, the MRF method produces a larger quantity of successful velocity measurements with errors comparable to those produced with the SRF method. However, this does not guarantee that the flows determined by either method will be correct. To ascertain the accuracy of each fitting method, we must compare the velocity measurements to a known velocity. We accomplish this by introducing known velocities into the data set through the tile tracking rate. Usually every tile is tracked longitudinally at a rate close to the surface differential rotation rate in order to minimize the effects of a large zonal flow. By altering the tracking rate for each tile such that it slides east or west at some Δv relative to the standard tracking rate, we introduce a known velocity as a Galilean transformation. We track the same set of tiles from the same temporal sequence at a variety of Δv values and use both fitting procedures to gather velocity measurements. By comparing the difference in measured zonal velocity between a region tracked at $\Delta v = 0$ and the same region tracked at a non-zero Δv , we determine how accurately each fitting method can reproduce the known velocity. By considering only the differences in measured velocities made at the same disk position, any systematics that are independent of tracking rate are removed

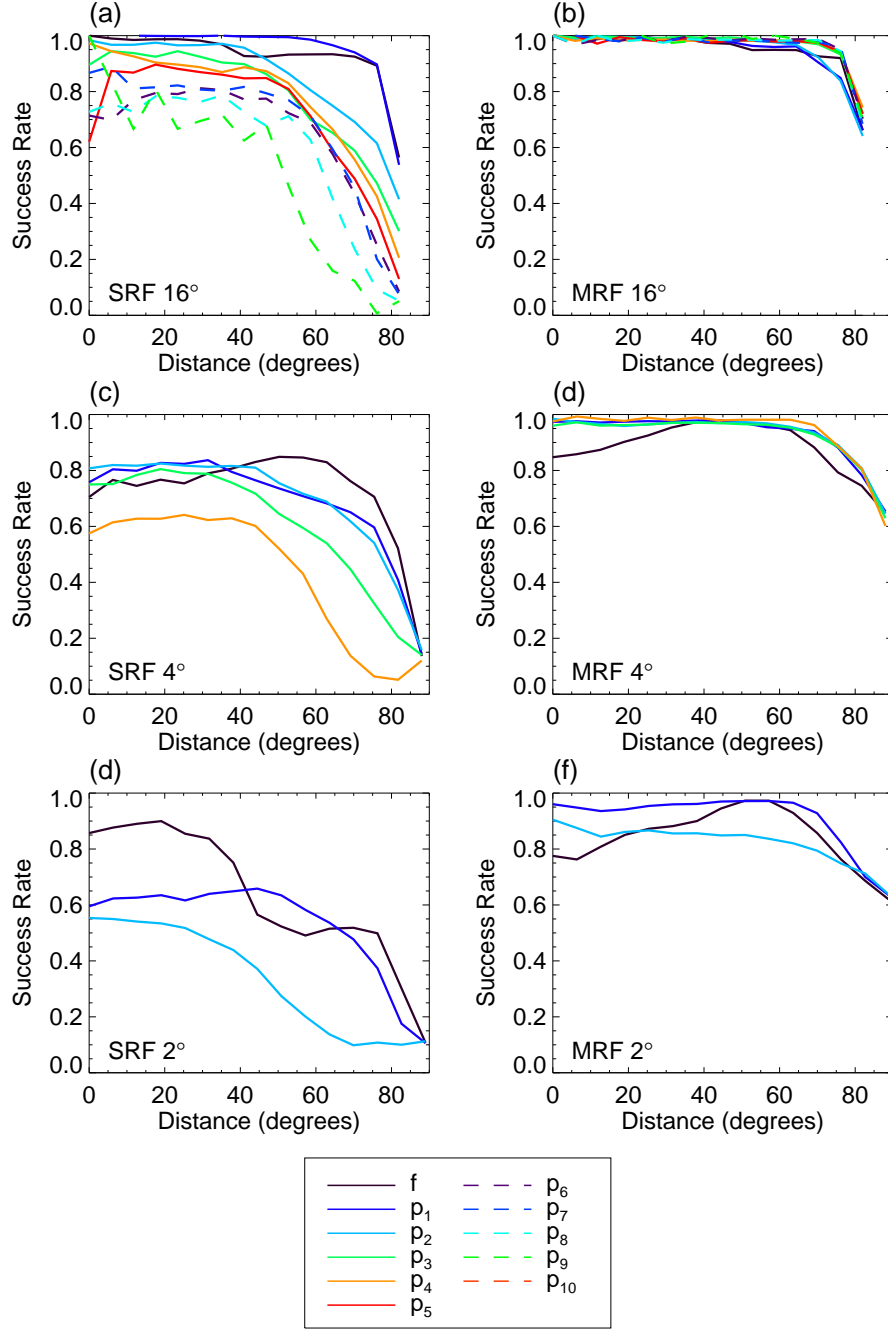


Figure 6: Success rate for each radial order averaged over all wavenumbers as a function of distance from disk center. The success rate as a function of distance from disk center for the SRF method is different for each radial order, while the success rate for the MRF method is largely constant. The f-mode has roughly the highest success rate, while each higher radial order has a successively lower success rate.

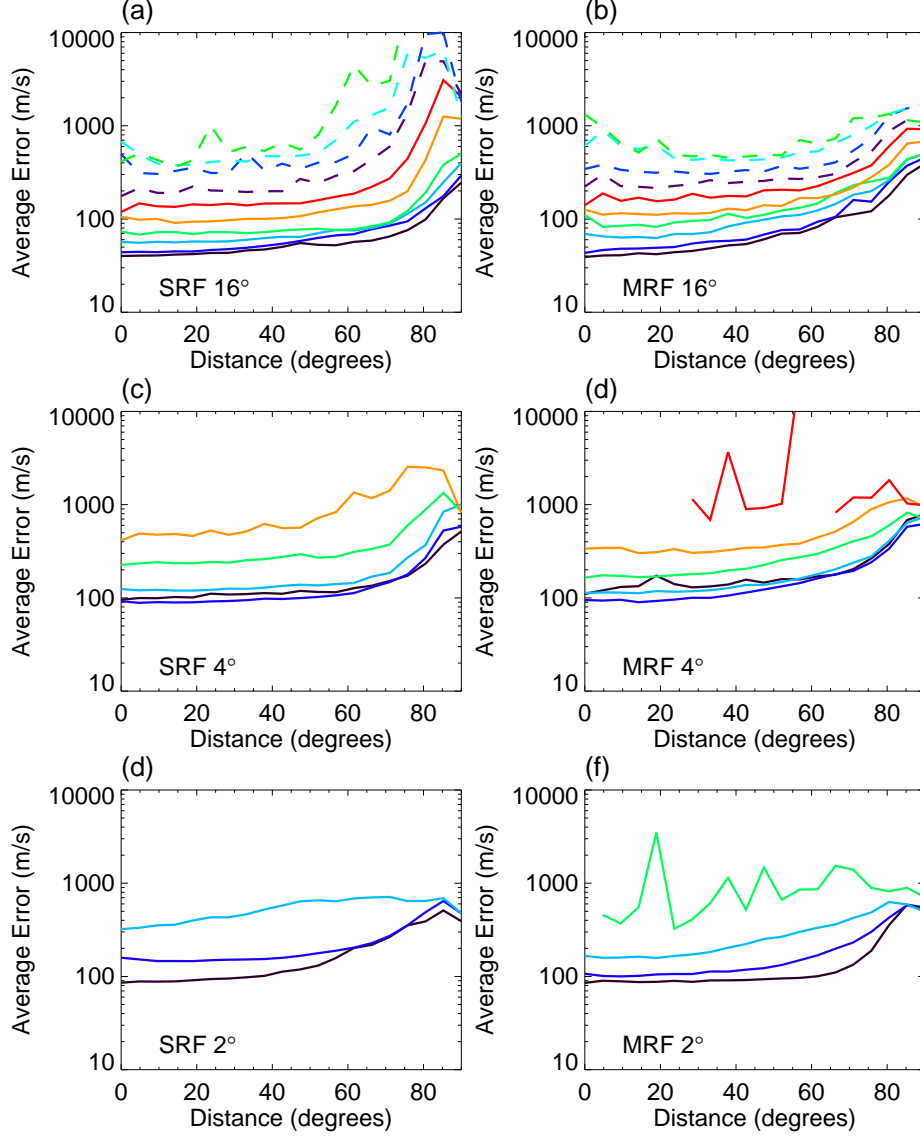


Figure 7: Average random error as a function of distance from disk center. The average error for both methods stay roughly constant near disk center and increase dramatically outside 60° . The f-mode has the lowest average error, while each higher radial order has a successively higher average error.

from the analysis.

A set of tiles covering the entire solar disk are tracked for a single day over a range of Δv spanning 500 m s^{-1} in each direction. The zonal-velocity measurements made at $\Delta v = 0$ are subtracted from those made at all other values of Δv independently for each mode of each tile to remove physical flows and systematics. The ideal velocity response is to have the measured velocity difference consistent with the introduced velocity difference. By subtracting the known velocity difference from the measured difference, we obtain the velocity deviation as a function of Δv . At each disk position, we perform a linear regression of these deviations against Δv for a collection of modes of similar phase-speed.

$$[u(\Delta v) - u_0] - \Delta v = a_0 + a_1 \Delta v \quad (8)$$

$$u_0 \equiv u(0) \quad (9)$$

Here, $u(\Delta v)$ is the measured velocity of a single mode obtained with either fitting procedure for a given tile tracked at Δv relative to the default rate, and u_0 is the measured velocity from the same tile tracked at the default rate. The left side of Equation 8 represents the deviation of measured velocity from the introduced velocity, while the right side expresses this deviation with a linear dependence on the introduced velocity. We bin together measurements made from modes within a range of phase-speeds at each disk position and fit the coefficients a_0 and a_1 . We use the coefficient a_1 as a diagnostic for how each method responds to velocities where $a_1 = 0$ is ideal. Mapped as a function of disk position (Figures 8, 9), this quantity shows the spatial uniformity of velocity response. The measured value of a_0 is consistent with zero in all cases, so it will not be discussed further. The mode set for each tile size is divided into two subsets based on the lower turning-point depth of each mode. The depth at which each mode set is split is dependent on tile size and is chosen to get roughly the same number of modes in each subset. Modes with lower turning points in the upper 16 Mm for 16° tiles, 4 Mm for 4° tiles, and 2 Mm for 2° tiles are shown in Figure 8, while modes with lower turning-point depths below this are shown in Figure 9.

The velocity response of modes with shallow lower turning points from 16° tiles is largely ideal across the entire disk for both the SRF and MRF methods (Figure 8), but with a slight overestimation of velocities at high latitudes. The MRF For smaller tile sizes, the SRF method produces position-dependent results, while the MRF method begins to show a uniform underestimation of velocities. For 4° tiles, the SRF method underestimates velocities far from the Equator by around 10 % and overestimates far from the central meridian by the same amount. The SRF pattern for 2° tiles is much less clear due to scatter in the measurements and lower success rate, but there is still a trend for overestimating in the east–west direction and underestimating in the north–south direction. The MRF results exhibit a much more spatially uniform pattern while still showing a consistent underestimation of 5 % to 10 %. For every tile size, the SRF results tend to be most accurate within 30° of disk center (the inner dashed circles of Figures 8 and 9) while the MRF results show a consistent underestimation that

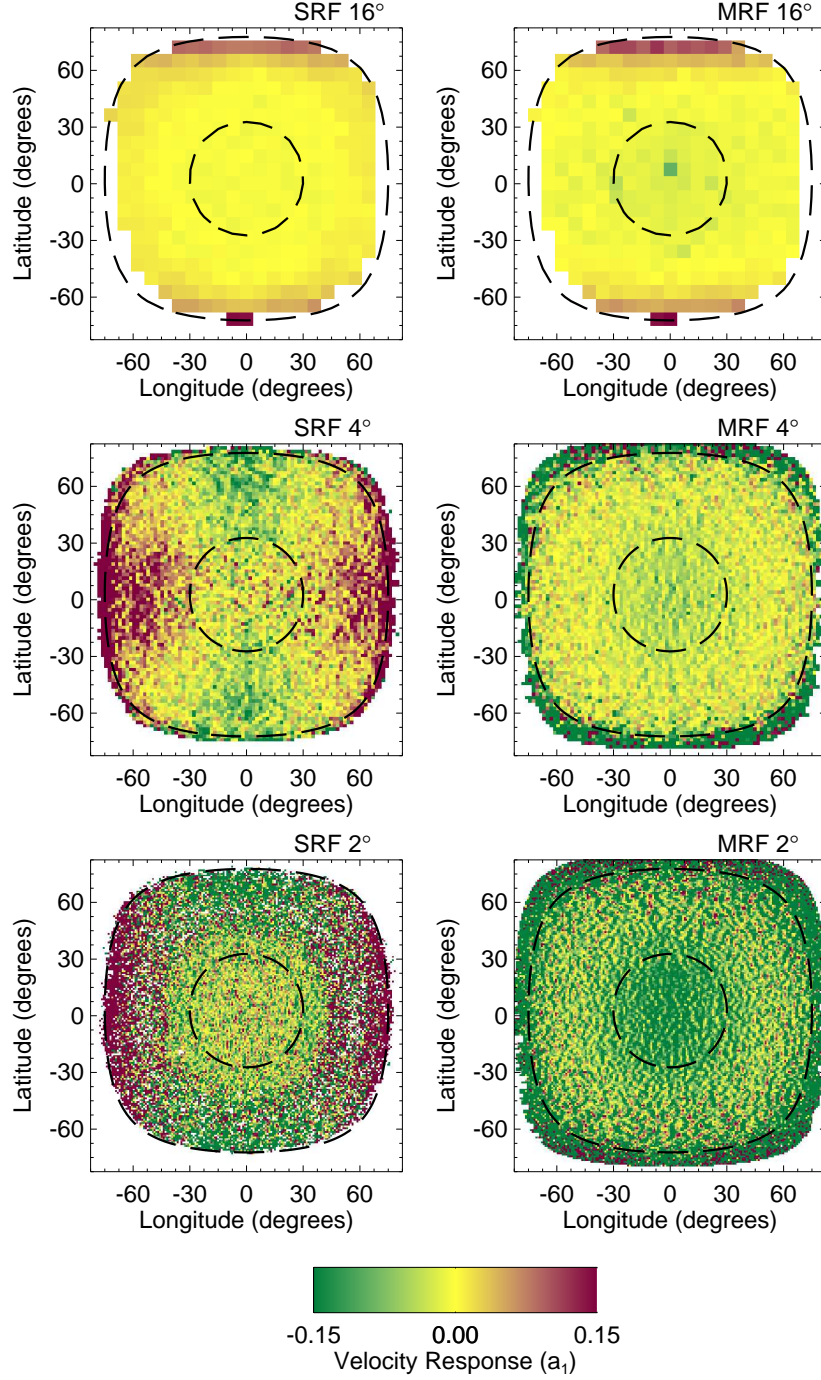


Figure 8: Velocity response for modes with lower turning points in the upper 16 Mm for 16° tiles, 4 Mm for 4° tiles, and 2 Mm for 2° tiles. The inner dashed circle is at 30° from disk center and the outer dashed line is at 75°. The SRF velocity response (left) has significantly higher spatial variability over the solar disk than the MRF method (right).

depends on tile size. From 30° out to 75° from disk center (the outer dashed lines of Figures 8 and 9), the MRF results are more accurate and spatially consistent than the SRF results. Beyond 75° for smaller tile sizes, the MRF results begin a consistent underestimation while the SRF results maintain the spatially varying pattern seen inside 75° . For modes sensitive to deeper flows (Figure 9), spatial uniformity is less of an issue while both methods obtain underestimations (5 % to 10 %) of the introduced velocity.

As tiles are tracked at large values of Δv , they begin to average over the sub-surface flows in a more extended physical region than the same tiles tracked at $\Delta v = 0$. Since this extra tracking distance is constant for all tile sizes, it is the smallest tile sizes that are affected the most. Sub-surface flows that vary on a length scale comparable to the tile size add to the scatter seen in Figures 8 and 9. Since the zonal component of these small-scale flows averages to zero over large regions of the disk, there is no significant bias introduced to the analysis. The small-scale pattern that emerges in the 2° map for MRF measurements (and 4° to a lesser extent) is due to supergranular flows. While the SRF maps should display the same exact pattern, it is interesting to note that the variance appears to obscure the pattern.

5 Discussion and Conclusions

Extracting frequency shifts that are interpreted as sub-surface velocities is done through fitting a model of acoustic power to a three-dimensional power spectrum. By fitting a model that includes various effects seen in the data to different segments of the three-dimensional spectrum, many unique fitting methods can be created. We have presented two such fitting procedures that differ both in the model used and the selection of data considered in a single optimization. These methods use identical expressions for the velocity-induced frequency splitting for each oscillation mode, yet produce significantly different results for the velocity measurements themselves.

5.1 Improvements in Depth for Inversions

The MRF method is able to obtain a greater number of velocity measurements for higher phase-speed modes in a given tile size than the SRF method. The increase in the success rate is also prominent for small tile sizes, suggesting that the treatment of mode blending plays an important role. Smaller tiles exhibit more leakage of modal power across wavenumbers than larger tiles due to spatial apodization. The spreading of power increases the effective width of modes at a single wavenumber. High-phase-speed modes are already spaced closer in wavenumber than other modes, enhancing the mode blending for their region of each spectrum. The SRF model does not accommodate any overlap of power from neighboring modes when fitting a single mode, causing a disparity between model and data that worsens at higher phase-speeds and smaller tile sizes. The MRF method, in accounting for mode overlap, returns a higher number of

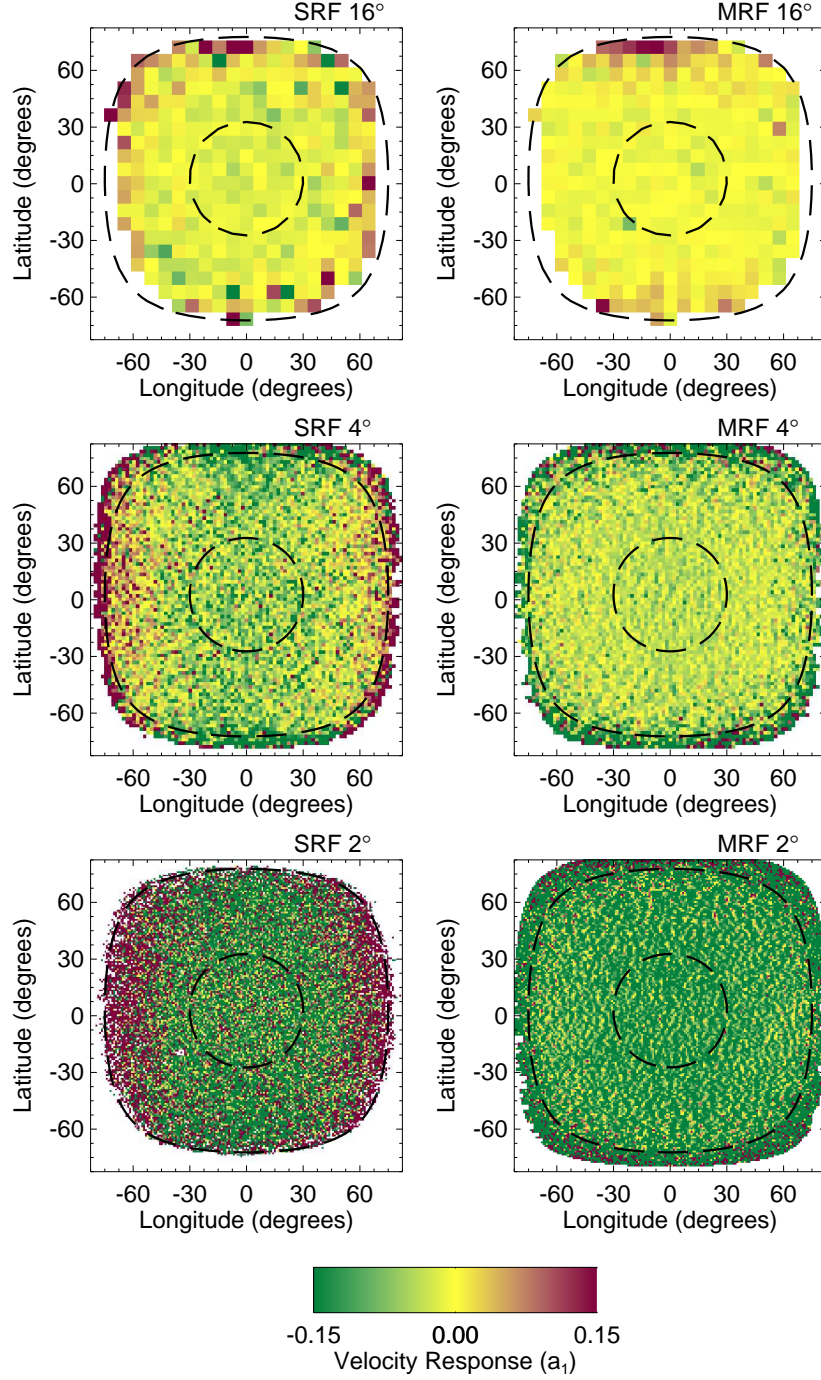


Figure 9: Velocity response for modes with lower turning points between 16 and 48 Mm for 16° tiles, between 4 and 12 Mm for 4° , and between 2 and 6 Mm for 2° . The inner dashed circle is at 30° from disk center and the outer dashed line is at 75° . Velocity measurements from deeper modes show less spatial variability over the solar disk for both the SRF method (left) and the MRF method (right).

successful measurements for these cases that can be used for studying deep flows and high horizontal resolution flows.

Referring to Equation 2, the set of measurements obtained from a single tile $\mathbf{u}_n(k)$ are measures of the true sub-surface velocity $\mathbf{v}(\mathbf{r})$ within a single region d^3r weighted with different sensitivity kernels $K(\mathbf{r})$. The sensitivity kernels generally do not have sufficient isolation of sensitivity at any given depth, so while having a larger number of high-phase-speed measurements results in reaching deeper into the Sun, these measurements alone do not provide a localized determination of sub-surface flows. The velocity measurements from many modes can be merged through an inversion to produce an estimate of the true sub-surface flow velocity in an isolated region. After finding a linear combination of sensitivity kernels that produces a single isolated peak of sensitivity at a target location, the estimated velocity for that region is constructed using the same linear combination of the associated velocity measurements:

$$\mathbf{w} = \sum_i a_i \mathbf{u}_i. \quad (10)$$

Here \mathbf{w} is the estimated sub-surface flow velocity at a selected target point, \mathbf{u}_i are the velocity measurements obtained from each mode (n, k) , and a_i are the coefficients of the linear combination determined by the inversion. To estimate the localization of each solution point, we can construct an averaging kernel that represents the spatial distribution of sensitivity for a solution point, just as a sensitivity kernel does for measured velocities:

$$H(\mathbf{r}) = \sum_i a_i K_i(\mathbf{r}). \quad (11)$$

The set of modes that we are able to fit reliably using either method determine the basis set of kernels that can be used to construct a locally isolated averaging kernel $H(\mathbf{r})$. It is not only how high in phase-speed the fitting method can reach, but also how reliably it can obtain successful measurements at all phase-speeds, that determines how well we can constrain sub-surface flows.

To determine how well averaging kernels can be constructed using the mode set available from each fitting method operating on 16° tiles, we use a simple 1D optimally localized averaging (OLA) inversion ([6]; [18]). The target functions are Gaussian profiles in depth and we restrict the set of sensitivity kernels available for inversion to those with a 80 % or higher success rate. Figure 10 compares averaging kernels constructed from each kernel set to the target functions used.

For shallow targets, both sets are able to accurately recreate the Gaussian target function with minimal sidelobes in depth. As the target depth increases, the number of kernels available for constructing averaging kernels decreases. Past a target depth of 30 Mm, the SRF set does not contain enough kernels to create a sufficiently isolated peak of sensitivity. The MRF set still contains enough unique kernels that reach past this depth to construct reasonable averaging kernels down to 45 Mm.

While the ability to construct reasonable averaging kernels through inversion is crucial for interpreting sub-surface flow results, this ignores the influence of measured

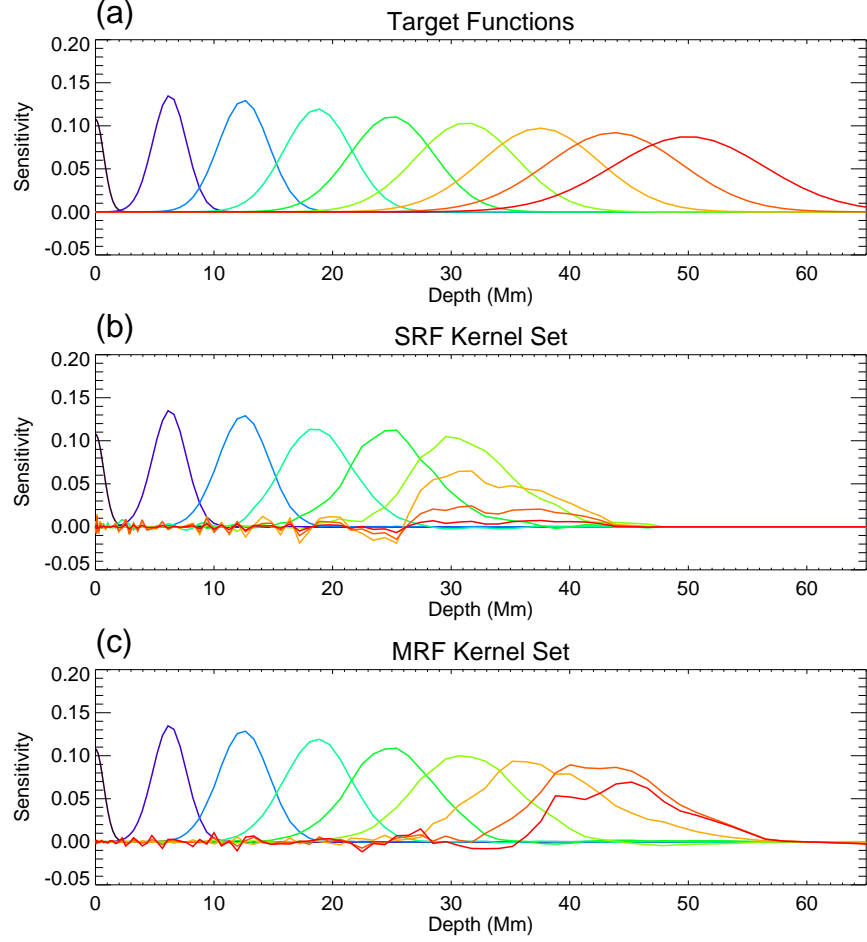


Figure 10: (a) Target functions for the 1D OLA inversion. (b) Averaging kernels constructed from sensitivity kernels from the SRF mode set. (c) Averaging kernels constructed from sensitivity kernels from the MRF mode set. The colors indicate different target depths. The SRF method is only able to produce averaging kernels down to a depth of about 30 Mm, while the MRF method reaches down to 45 Mm.

errors. For OLA inversions, an optimization is done to balance the shape of the averaging kernel against the propagated error:

$$\sigma^2 = \sum_i (a_i \sigma_i)^2. \quad (12)$$

In this way, it is both the available mode set and the associated errors that determine the usefulness of the inversion results. Figure 11 shows the propagated inversion error as a function of the averaging kernel center-of-gravity depth for each fitting method. To mimic the result of inverting an average of many days of data over all longitudes to look at large-scale mean flows, the error for each mode has been scaled:

$$\sigma_i \rightarrow \frac{\sigma_i}{\sqrt{N_d N_l r_i}}, \quad (13)$$

where r_i is the success rate for a given mode, N_d is the number of days to average over, and N_l is the number of longitudes to average over. We set N_d to 180 to mimic inversions of six months of data and N_l to 17 for the typical number of tiles spanning the solar disk in longitude. The various curves for different regions of the solar disk demonstrate how much the dependence of the success rate and average error on disk position affect inversion results. The sharp downturn seen at the deep end of each curve indicates where the averaging kernels no longer resemble the target functions and are therefore useless.

Due to the higher success rate for each mode, the MRF method is able to produce useful inversion results to a greater depth than the SRF method. As the distance from disk center increases, the MRF results remain mostly constant while the SRF results start to deteriorate both in error magnitude and kernel depth. Despite larger average error in the MRF results for shallow modes seen in Figure 5, the higher success rate for these modes causes the inversion errors to not show a significant increase relative to the SRF results.

5.2 Spatial Variability and Implications

Anisotropy along the azimuthal direction of the power spectrum appears to be the primary cause of fitting problems near the solar limb. The SRF method tackles this issue by operating on spectra that have been “flattened” along azimuth. This step not only amplifies the modal power in certain directions in order to eliminate anisotropy, but also amplifies the noise. By accounting for the natural variation of power along each ridge as well as in the background power, the MRF method is able to push closer to the solar limb. As seen in Figure 6, the mode set measureable by the SRF method is position-dependent. The higher-order modes see a drop in success rate closer to disk center than lower-order modes. The depth to which one can measure is then shallower near the limb. The MRF method, in contrast, allows for consistent determinations of sub-surface flows within 75° of disk center.

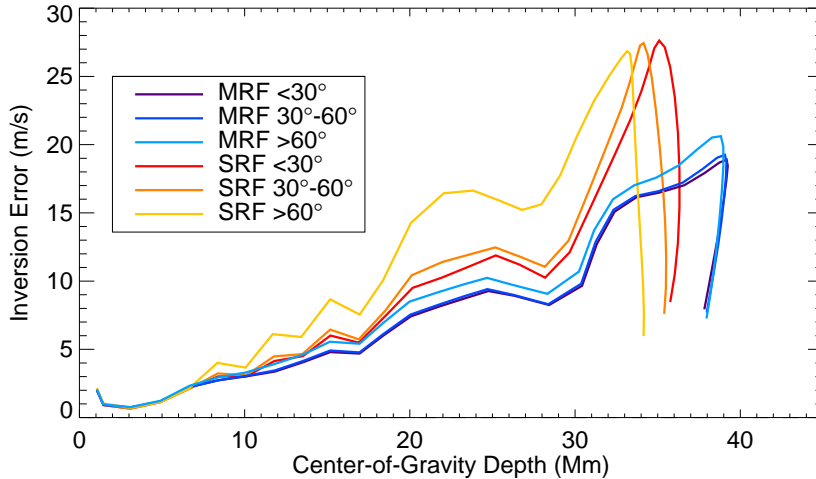


Figure 11: Typical error magnitudes propagated through 1D OLA inversion as a function of depth for multiple distances from disk center. The sharp turnoff at the deep end of each curve demonstrates where proper averaging kernels can no longer be produced. The inversions done with SRF results tend to turn off between 33 and 35 Mm depth while those done with MRF results turn off between 38 and 40 Mm.

Large tile sizes tend to provide the most accurate velocity measurements for both fitting methods, with a slight tendency for both to overestimate at high latitudes (Figure 8). Both methods exhibit spatial variability in the velocity response for smaller tile sizes, but only for low-phase-speed modes. The MRF response is radially symmetric with the worst velocity response occurring near disk center, while the worst regions of the SRF response appear far from disk center and have opposite signs between the north–south and east–west direction. While the magnitude of this variability is relatively small, there are significant implications when attempting to analyze small residual flows left over from the subtraction of large-scale flows. Tiles at high latitudes tracked at the Carrington rate see an overall flow speed of nearly 200 m s^{-1} due to differential rotation. The systematic underestimation seen in the SRF method introduces an anomalous retrograde flow of around 20 m s^{-1} , similar to the magnitude of the center-to-limb velocity systematic for low-phase-speed modes [8]. It is unclear what causes the systematic inaccuracy in either fitting method, although the strong dependence on disk position suggests a possible coupling between velocity measurements and power anisotropy.

As mentioned in Section 3.2, there is a significant increase in the computational cost when switching from the SRF method to the MRF method. While processing a day’s worth of tiles with the SRF method is typically one of the fastest steps in the HMI Pipeline, the MRF method brings the process of fitting more in line with the time needed for tracking. It is important to consider when this additional computational burden is justified. The SRF method produces high-quality results for low-phase-speed

modes and for large tiles near disk center. This leaves three distinct cases where the MRF method provides improvements: measuring higher-phase-speed modes, pushing closer to the limb, and using small tiles. The implications of this new procedure on determining sub-surface flows follow these three technical improvements. For a given tile size, we are able to extend our analysis deeper into the Sun while maintaining a constant horizontal and temporal resolution. Analysis to these extended depths can be performed consistently across most of the solar disk, providing uniform coverage over a larger fraction of the Sun. The increased reliability of small tile sizes permits higher horizontal resolution analysis of sub-surface flows.

Acknowledgements This work was supported by NASA through NASA grants NNX08AJ08G, NNX08AQ28G, and NNX09AB04G. The data used here are courtesy of NASA/SDO and the HMI science team. SDO is a NASA mission, and the HMI project is supported by NASA contract NAS5-02139.

References

- [1] E. R. Anderson, T. L. Duvall, Jr., and S. M. Jefferies. Modeling of solar oscillation power spectra. *Astrophys. J.*, 364:699–705, December 1990.
- [2] S. Basu, H. M. Antia, and S. C. Tripathy. Ring diagram analysis of near-surface flows in the sun. *Astrophys. J.*, 512:458–470, February 1999.
- [3] A. C. Birch, L. Gizon, B. W. Hindman, and D. A. Haber. The linear sensitivity of helioseismic ring diagrams to local flows. *Astrophys. J.*, 662:730–737, June 2007.
- [4] R. S. Bogart, C. Baldner, S. Basu, D. A. Haber, and M. C. Rabello-Soares. HMI ring diagram analysis I. The processing pipeline. *J. Phys. Conf. Ser.*, 271(1):012008, January 2011.
- [5] J. Christensen-Dalsgaard, W. Dappen, S. V. Ajukov, E. R. Anderson, H. M. Antia, S. Basu, and et al. The current state of solar modeling. *Science*, 272:1286–1292, May 1996.
- [6] J. Christensen-Dalsgaard, J. Schou, and M. J. Thompson. A comparison of methods for inverting helioseismic data. *Mon. Not. Roy. Astron. Soc.*, 242:353–369, February 1990.
- [7] I. González Hernández, R. Howe, R. Komm, and F. Hill. Meridional circulation during the extended solar minimum: Another component of the torsional oscillation? *Astrophys. J. Lett.*, 713:L16–L20, April 2010.
- [8] B. Greer, B. Hindman, and J. Toomre. Center-to-Limb Velocity Systematic in Ring-Diagram Analysis. In K. Jain, S. C. Tripathy, F. Hill, J. W. Leibacher, and A. A. Pevtsov, editors, *Astronomical Society of the Pacific Conference Series*,

volume 478 of *Astronomical Society of the Pacific Conference Series*, page 199, December 2013.

- [9] D. A. Haber, B. W. Hindman, J. Toomre, R. S. Bogart, R. M. Larsen, and F. Hill. Evolving submerged meridional circulation cells within the upper convection zone revealed by ring-diagram analysis. *Astrophys. J.*, 570:855–864, May 2002.
- [10] D. A. Haber, B. W. Hindman, J. Toomre, R. S. Bogart, M. J. Thompson, and F. Hill. Solar shear flows deduced from helioseismic dense-pack samplings of ring diagrams. *Solar Phys.*, 192:335–350, March 2000.
- [11] J. Harvey. High-resolution helioseismology. In E. Rolfe and B. Battrock, editors, *Future missions in solar, heliospheric & space plasma physics*, volume 235 of *ESA Special Publication*, pages 199–208, June 1985.
- [12] F. Hill. Rings and trumpets - Three-dimensional power spectra of solar oscillations. *Astrophys. J.*, 333:996–1013, October 1988.
- [13] B. W. Hindman, D. Gough, M. J. Thompson, and J. Toomre. Helioseismic ring analyses of artificial data computed for two-dimensional shearing flows. *Astrophys. J.*, 621:512–523, March 2005.
- [14] R. Komm, I. González Hernández, F. Hill, R. Bogart, M. C. Rabello-Soares, and D. Haber. Subsurface meridional flow from HMI using the ring-diagram pipeline. *Solar Phys.*, 287:85–106, October 2013.
- [15] L. Rightmire-Upton, D. H. Hathaway, and K. Kosak. Measurements of the sun’s high-latitude meridional circulation. *Astrophys. J. Lett.*, 761:L14, December 2012.
- [16] S. Routh, D. A. Haber, B. W. Hindman, R. S. Bogart, and J. Toomre. The influence of tracking rate on helioseismic flow inferences. *J. Phys. Conf. Ser.*, 271(1):012014, January 2011.
- [17] P. H. Scherrer, J. Schou, R. I. Bush, A. G. Kosovichev, R. S. Bogart, J. T. Hoeksema, Y. Liu, T. L. Duvall, J. Zhao, A. M. Title, C. J. Schrijver, T. D. Tarbell, and S. Tomczyk. The Helioseismic and Magnetic Imager (HMI) investigation for the Solar Dynamics Observatory (SDO). *Solar Phys.*, 275:207–227, January 2012.
- [18] J. Schou, H. M. Antia, S. Basu, R. S. Bogart, R. I. Bush, S. M. Chitre, and et al. Helioseismic studies of differential rotation in the solar envelope by the solar oscillations investigation using the Michelson Doppler Imager. *Astrophys. J.*, 505:390–417, September 1998.
- [19] H. B. Snodgrass. Separation of large-scale photospheric Doppler patterns. *Solar Phys.*, 94:13–31, August 1984.

- [20] J. Zhao, R. S. Bogart, A. G. Kosovichev, T. L. Duvall, Jr., and T. Hartlep. Detection of equatorward meridional flow and evidence of double-cell meridional circulation inside the sun. *Astrophys. J. Lett.*, 774:L29, September 2013.
- [21] J. Zhao, K. Nagashima, R. S. Bogart, A. G. Kosovichev, and T. L. Duvall, Jr. Systematic center-to-limb variation in measured helioseismic travel times and its effect on inferences of solar interior meridional flows. *Astrophys. J. Lett.*, 749:L5, April 2012.